arXiv:1407.3512v1 [cs.AI] 13 Jul 2014

A New Rational Algorithm for View Updating in Relational Databases *

Radhakrishnan Delhibabu[1] and Andreas Behrend[2]

[1] Informatik 5, Knowledge-Based Systems Group
RWTH Aachen, Germany
delhibabu@kbsg.rwth-aachen.de
[2] Informatik 3, IDB Group
University of Bonn, Germany
behrend@cs.uni-bonn.de

Abstract. The dynamics of belief and knowledge is one of the major components of any autonomous system that should be able to incorporate new pieces of information. In order to apply the rationality result of belief dynamics theory to various practical problems, it should be generalized in two respects: first it should allow a certain part of belief to be declared as immutable; and second, the belief state need not be deductively closed. Such a generalization of belief dynamics, referred to as base dynamics, is presented in this paper, along with the concept of a generalized revision algorithm for knowledge bases (Horn or Horn logic with stratified negation). We show that knowledge base dynamics has an interesting connection with kernel change via hitting set and abduction. In this paper, we show how techniques from disjunctive logic programming can be used for efficient (deductive) database updates. The key idea is to transform the given database together with the update request into a disjunctive (datalog) logic program and apply disjunctive techniques (such as minimal model reasoning) to solve the original update problem. The approach extends and integrates standard techniques for efficient query answering and integrity checking. The generation of a hitting set is carried out through a hyper tableaux calculus and magic set that is focused on the goal of minimality.

Keyword: AGM, Belief Revision, Knowledge Base Dynamics, Kernel Change, Abduction, Hyber Tableaux, Magic Set, View update, Update Propagation.

1 Introduction

Modeling intelligent agents' reasoning requires designing knowledge bases for the purpose of performing symbolic reasoning. Among the different types of knowledge representations in the domain of artificial intelligence, logical representations stem from classical logic. However, this is not suitable for representing

* This paper extends work from Behrend [6] and Delhibabu [18].

or treating items of information containing vagueness, incompleteness or uncertainty, or knowledge base evolution that leads the agent to change his beliefs about the world.

When a new item of information is added to a knowledge base, it may become inconsistent. Revision means modifying the knowledge base in order to maintain consistency, while keeping the new information and removing (contraction) or not removing the least possible previous information. In our case, update means revision and contraction, that is insertion and deletion from a database perspective. Previous work [18] and [10,11] makes connections with revision from knowledge base dynamics.

Our knowledge base dynamics is defined in two parts: an immutable part (formulae) and updatable part (literals) (for definition and properties see works of Nebel [46] and Segerberg [53]). Knowledge bases have a set of integrity constraints. In the case of finite knowledge bases, it is sometimes hard to see how the update relations should be modified to accomplish certain knowledge base updates.

Example 1. Consider a database with an (immutable) rule that a staff member is a person who is currently working in a research group under a chair. Additional (updatable) facts are that matthias and gerhard are group chairs, and delhibabu and aravindan are staff members in group info1. Our first integrity constraint (IC) is that each research group has only one chair i.e., $\forall x, y, z$ (y=z) $\leftarrow$ group_chair(x,y) $\wedge$ group_chair(x,z). Second integrity constraint is that a person can be a chair for only one research group, i.e., $\forall x, y, z$ (y=z)$\leftarrow$ group_chair(y,x) $\wedge$ group_chair(z,x).

Immutable part: staff_chair(X,Y)$\leftarrow$ staff_group(X,Z),group_chair(Z,Y).

Updatable part: group_chair(infor1,matthias)$\leftarrow$
group_chair(infor2,gerhard)$\leftarrow$
staff_group(delhibabu,infor1)$\leftarrow$
staff_group(aravindan,infor1)$\leftarrow$

Suppose we want to update this database with the information, staff_chair(aravindan,gerhard); From the immutable part, we can deduce that this can be achieved by asserting staff_group(aravindan,Z) $\bigwedge$ group_chair(Z,gerhard)

If we are restricted to definite clauses, there are three plausible ways to do this. When dealing with the revision of a knowledge base (both insertions and deletions), there are other ways to change a knowledge base and it has to be performed automatically too. Considering the information, change is precious and must be preserved as much as possible. The *principle of minimal change* [25,52] can provide a reasonable strategy. On the other hand, practical implementations have to handle contradictory, uncertain, or imprecise information, so several problems can arise: how to define efficient change in the style of Carlos Alchourrón, Peter Gärdenfors, and David Makinson (AGM) [1]; what result has to be chosen [30,38,44]; and finally, according to a practical point of view, what computational model to support for knowledge base revision has to be provided?

The basic idea in [6,12] is to employ the model generation property of hyper tableaux and magic set to generate models, and read off diagnosis from them. One specific feature of this diagnosis algorithm is the use of semantics (by transforming the system description and the observation using an initial model of the correctly working system) in guiding the search for a diagnosis. This semantical guidance by program transformation turns out to be useful for database updates as well. More specifically we use a (least) Herbrand model of the given database to transform it along with the update request into a disjunctive logic program in such a way that the models of this transformed program stand for possible updates.

We discuss two ways of transforming the given database together with the view update (insert and delete) request into a disjunctive logic program resulting in two variants of view update algorithms. In the first variant, a simple and straightforward transformation is employed. Unfortunately, not all models of the transformed program represent a rational update using this approach. The second variant of the algorithm uses the least Herbrand model of the given database for the transformation. In fact what we referred to as offline preprocessing before is exactly this computation of the least Herbrand model. This variant is very meaningful in applications where views are materialized for efficient query answering. The advantage of using the least Herbrand model for the transformation is that all models of the transformed disjunctive logic program (not just the minimal ones) stand for a rational update.

The rest of paper is organized as follows: First we start with preliminaries in Section 2. In Section 3, we introduce knowledge base dynamics along with the concept of generalized revision, and revision operator for knowledge base. Section 4 studies the relationship between knowledge base dynamics and abduction. We discuss an important application of knowledge base dynamics in providing an axiomatic characterization for updating view literals over databases. We briefly discuss hyper tableaux calculus and magic set in Section 5. We present two variants of our rational and efficient algorithm for view updating in Section 6. In Section 7, we give a brief overview of related work. In Section 8 we draw conclusions with a summary of our contribution and indicate future directions of our investigation. All proofs can be found in the Appendix.

2 Background

2.1 Rationality of change

Rationality of change has been studied at an abstract philosophical level by various researchers, resulting in well known AGM Postulates for revision [1,23]. However, it is not clear how these rationality postulates can be applied in real world problems such as database updates and this issue has been studied in detail by works such as [18]. In the sequel, we briefly recall the postulates and an algorithm for revision based on abduction from Delhibabu [18] work.

We consider a propositional language $\mathcal{L}_{\mathcal{P}}$ defined from a finite set of propositional variables $\mathcal{P}$ and the standard connectives. We use lower case Roman

letters $a, b, x, y, ...$ to range over elementary letters and Greek letters $\varphi, \phi, \psi, ...$ for propositional formulae. Sets of formulae are denoted by upper case Roman letters $A, B, F, K, ...$ A literal is an atom (positive literal), or a negation of an atom (negative literal).

Formally, a finite Horn knowledge base KB (Horn [17] or Horn logic with stratified negation [26]) is defined as a finite set of formulae from language $\mathcal{L}_{\mathcal{H}}$, and divided into three parts: an immutable theory KB_I is a Horn formula, which is the fixed part of the knowledge; updatable theory KB_U is a Horn clause; and integrity constraint KB_{IC} representing a set of clauses (Horn logic with stratified negation).

Definition 1 (Horn Knowledge Base). *A Horn knowledge base, KB is a finite set of Horn formulae from language $\mathcal{L}_{\mathcal{H}}$, s.t $KB = KB_I \cup KB_U \cup KB_{IC}$, $KB_I \cap KB_U = \varnothing$ and $KB_U \cap KB_{IC} = \varnothing$.*

In the AGM approach, a belief is represented by a sentence over a suitable language $\mathcal{L}_{\mathcal{H}}$, and a belief KB is represented by a set of sentence that are close wrt the logical closure operator Cn. It is assumed that $\mathcal{L}_{\mathcal{H}}$, is closed under application of the boolean operators negation, conjunction, disjunction, and implication.

Definition 2. *Let KB be a knowledge base with an immutable part KB_I. Let α and β be any two (Horn or Horn logic with stratified negation) clauses from $\mathcal{L}_{\mathcal{H}}$. Then, α and β are said to be* KB-equivalent *iff the following condition is satisfied: $\forall$ set of Horn clauses $E \subseteq \mathcal{L}_{\mathcal{H}}$: $KB_I \cup E \vdash \alpha$ iff $KB_I \cup E \vdash \beta$.*

The revision can be trivially achieved by expansion, and the axiomatic characterization could be straightforwardly obtained from the corresponding characterizations of the traditional models [21]. The aim of our work is not to define revision from contraction, but rather to construct and axiomatically characterize revision operators in a direct way.

These postulates stem from three main principles: the new item of information has to appear in the revised knowledge base, the revised base has to be consistent and revision operation has to change the least possible beliefs. Now we consider the revision of a Horn (Horn logic with stratified negation) clause α wrt KB, written as $KB * \alpha$. The rationality postulates for revising α from KB can be formulated as follows:

Definition 3 (Rationality postulates for knowledge base revision).

*(KB*1)* Closure: *$KB * \alpha$ is a knowledge base.*
*(KB*2)* Weak Success: *if α is consistent with $KB_I \cup KB_{IC}$ then $\alpha \subseteq KB * \alpha$.*
*(KB*3.1)* Inclusion: *$KB * \alpha \subseteq Cn(KB \cup \alpha)$.*
*(KB*3.2)* Immutable-inclusion: *$KB_I \subseteq Cn(KB * \alpha)$.*
*(KB*4.1)* Vacuity 1: *if α is inconsistent with $KB_I \cup KB_{IC}$ then $KB * \alpha = KB$.*
*(KB*4.2)* Vacuity 2: *if $KB \cup \alpha \nvdash \bot$ then $KB * \alpha = KB \cup \alpha$.*
*(KB*5)* Consistency: *if α is consistent with $KB_I \cup KB_{IC}$ then $KB * \alpha$ is consistent with $KB_I \cup KB_{IC}$.*

*(KB*6)* Preservation: *If α and β are KB-equivalent, then $KB * \alpha \leftrightarrow KB * \beta$.*

*(KB*7.1)* Strong relevance: *$KB * \alpha \vdash \alpha$ If $KB_I \nvdash \neg\alpha$*

*(KB*7.2)* Relevance: *If $\beta \in KB \backslash KB * \alpha$, then there is a set KB' such that $KB * \alpha \subseteq KB' \subseteq KB \cup \alpha$, KB' is consistent $KB_I \cup KB_{IC}$ with α, but $KB' \cup \{\beta\}$ is inconsistent $KB_I \cup KB_{IC}$ with α.*

*(KB*7.3)* Weak relevance: *If $\beta \in KB \backslash KB * \alpha$, then there is a set KB' such that $KB' \subseteq KB \cup \alpha$, KB' is consistent $KB_I \cup KB_{IC}$ with α, but $KB' \cup \{\beta\}$ is inconsistent $KB_I \cup KB_{IC}$ with α.*

Now we recall an algorithm for revision based on abduction presented in [2,18], Some basic definitions required for the algorithm are presented first.

Definition 4 (Minimal abductive explanation). *Let KB be a Horn knowledge base and α an observation to be explained. Then, for a set of abducibles (KB^I), Δ is said to be an abductive explanation wrt KB^I iff $KB^I \cup \Delta \vdash \alpha$. Δ is said to be* minimal*[51] wrt KB^I iff no proper subset of Δ is an abductive explanation for α, i.e., $\nexists \Delta'$ s.t. $KB^I \cup \Delta' \vdash \alpha$.*

Definition 5 (Local minimal abductive explanations). *Let $KB^{I'}$ be a smallest subset of KB^I, s.t Δ is a minimal abductive explanation of α wrt $KB^{I'}$ (for some Δ). Then Δ is called local minimal [15,39] for α wrt KB^I.*

The general revision algorithm of [18] is reproduced here as Algorithm 1. The basic idea behind this algorithm is to generate all (locally minimal) explanations for the sentence to be contracted and determine a hitting set for these explanations. Since all (locally minimal) explanations are generated this algorithm is of exponential space and time complexity

Definition 6 (Hitting set). *Let S be a set of sets. Then a set HS is a* hitting set *of S iff $HS \subseteq \cup S$ and for every non-empty element R of S, $R \cap HS$ is non empty.*

Algorithm 1	Generalized revision algorithm
Input :	A Horn knowledge base $KB = KB_I \cup KB_U \cup KB_{IC}$ and a Horn clause α to be revised.
Output:	A new Horn knowledge base $KB' = KB_I \cup KB^*_U \cup KB_{IC}$, s.t. KB'is a generalized revision α to KB.
Procedure $KB(KB, \alpha)$	
begin	
1.	Let V:= $\{c \in KB_{IC} \mid KB_I \cup KB_{IC}$ inconsistent with α wrt $c\}$ $P := N := \emptyset$ and $KB' = KB$
2.	While $(V \neq \emptyset)$ select a subset $V' \subseteq V$ For each $v \in V'$, select a literal to be remove (add to N) or a literal to be added(add to P) Let KB $:= KR(KB, P, N)$ Let V:= $\{c \in KB_{IC} \mid KB_I$ inconsistent with α wrt $c\}$ return
3.	Produce a new Horn knowledge base KB'
end.	

Algorithm 2

Procedure $KR(KB, \Delta^+, \Delta^-)$

begin

1. Let $P := \{e \in \Delta^+ \mid KB_I \not\models e\}$ and $N := \{e \in \Delta^- \mid KB_I \models e\}$

2. While $(P \neq \emptyset)$ or $(N \neq \emptyset)$

select a subset $P' \subseteq P$ *or* $N' \subseteq N$

Construct a set $S_1 = \{X \mid X$ is a KB-closed locally minimal abductive wrt P explanation for α wrt $KB_I\}$.

Construct a set $S_2 = \{X \mid X$ is a KB-closed locally minimal abductive wrt N explanation for α wrt $KB_I\}$.

Determine a hitting set $\sigma(S_1)$ and $\sigma(S_2)$

If $((N' = \emptyset)$ *and* $(P' \neq \emptyset))$

Produce $KB' = KB_I \cup \{(KB_U \cup \sigma(S_1)\}$

else

Produce $KB' = KB_I \cup \{(KB_U \backslash \sigma(S_2) \cup \sigma(S_1)\}$

end if

If $((N' \neq \emptyset)$ and $(P' = \emptyset))$

Produce $KB' = KB_I \cup \{(KB_U \backslash \sigma(S_2)\}$

else

Produce $KB' = KB_I \cup \{(KB_U \backslash \sigma(S_2) \cup \sigma(S_1)\}$

end if

Let $P := \{e \in \Delta^+ \mid KB_I \not\models e\}$ and $N := \{e \in \Delta^- \mid KB_I \models e\}$

3. return KB'

end.

Theorem 1. *Let KB be a knowledge base and α is (Horn or Horn logic with stratified negation) formula.*

1. *If Algorithm 1 produced KB' as a result of revising α from KB, then KB' satisfies all the rationality postulates (KB*1) to (KB*6) and (KB*7.3).*
2. *Suppose KB'' satisfies all these rationality postulates for revising α from KB, then KB'' can be produced by Algorithm 1.*

3 Deductive database

A *Deductive database* DDB consists of three parts: an *intensional database* IDB (KB_I), a set of definite program clauses, *extensional database* EDB (KB_U), a set of ground facts; and *integrity constraints* IC. The intuitive meaning of DDB is provided by the *Least Herbrand model semantics* and all the inferences are carried out through *SLD-derivation. All the predicates that are defined in* IDB *are referred to as* view predicates *and those defined in* EDB *are referred to as* base predicates*.* Extending this notion, an atom with a view predicate is said to be a *view atom,*and similarly an atom with base predicate is a *base atom.* Further we assume that IDB does not contain any unit clauses and no predicate defined in a given DDB is both view and base.

Two kinds of view updates can be carried out on a DDB: An atom, that does not currently follow from DDB, can be *inserted,* or an atom, that currently follows from DDB can be *deleted.* When an atom A is to be updated, the view update problem is to insert or delete only some relevant EDB facts, so that the modified EDB together with IDB will satisfy the updating of A to DDB.

Note that a DDB can be considered as a knowledge base to be revised. The IDB is the immutable part of the knowledge base, while the EDB forms the updatable part. In general, it is assumed that the language underlying a DDB is fixed and the semantics of DDB is the least Herbrand model over this fixed language. We assume that there are no function symbols implying that the Herbrand Base is finite. Therefore, the IDB is practically a shorthand of its ground instantiation[3] written as IDB_G. In the sequel, technically we mean IDB_G when we refer simply to IDB. Thus, a DDB represents a knowledge base where the immutable part is given by IDB_G and updatable part is the EDB. Hence, the rationality postulates (KB*1)-(KB*6) and (KB*7.3) provide an axiomatic characterization for update (insert and delete) a view atom A from a definite database DDB.

Logic provides a conceptual level for understanding the meaning of relational databases. Hence, the rationality postulates (KB*1)-(KB*6) and (KB*7.3) can provide an axiomatic characterization for view updates in relational databases too. A relational database together with its view definitions can be represented by a deductive database (EDB representing tuples in the database and IDB representing the view definitions), and so the same algorithm can be used to delete view extensions from relational deductive databases.

[3] a ground instantiation of a definite program P is the set of clauses obtained by substituting terms in the Herbrand Universe for variables in P in all possible ways

3.1 Disjunctive Deductive Databases

A disjunctive Datalog rule is a function-free clause of the form $H_1 \vee \ldots \vee H_m \leftarrow L_1 \wedge \ldots \wedge L_n$ with $m, n \geq 1$ where the rule's head $H_1 \vee \ldots \vee H_m$ is a disjunction of positive atoms, and the rule's body $L_1 \wedge \ldots \wedge L_n$ consists of literals, i.e., positive or negative atoms, if only positive atoms then (definite deductive) database. If $H \equiv H_1 \vee \ldots \vee H_m$ is the head of a given rule IDB, we use $pred(IDB)$ to refer to the set of predicate symbols of H, i.e., $pred(IDB) = \{pred(H_1), \ldots, pred(H_m)\}$. For a set of rules IDB, $pred(IDB)$ is defined again as $\bigcup_{r \in IDB} pred(r)$. A *disjunctive fact* $f \equiv f_1 \vee \ldots \vee f_k$ is a disjunction of ground atoms f_i with $i \geq 1$. f is called *definite* if $i = 1$. In the following, we identify a disjunctive fact with a set of atoms such that the occurrence of a ground atom A within a fact f can also be written as $A \in f$. The set difference operator can then be used to exclude certain atoms from a disjunction while the empty set is interpreted as the boolean constant false.

A disjunctive deductive database $DDDB$ is a pair $\langle IDB, EDB, IC \rangle$ where EDB is a finite set of disjunctive facts and IDB a finite set of disjunctive rules such that $pred(EDB) \cap pred(IDB) = \emptyset$. Again, stratifiable (definite) deductive rules are considered only, that is, recursion through negative predicate occurrences is not permitted. In addition to the usual stratification concept for definite rules it is required that all predicates within a rule's head are assigned to the same stratum.

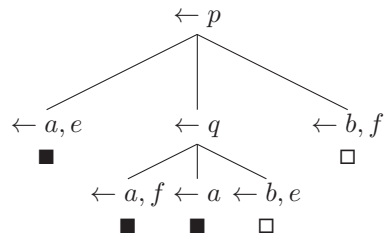
An update request U = B, where B is a set of base facts, is not true in KB. Then, we need to find a transaction $T = T_{ins} \cup T_{del}$, where $T_{ins}(\Delta_i)$ (resp. $T_{del}(\Delta_j)$) is the set of facts, such that U is true in $DDB' = ((EDB - T_{del} \cup T_{ins}) \cup IDB \cup IC)$. Since we consider stratifiable (definite) deductive databases, SLD-trees can be used to compute the required abductive explanations. The idea is to get all EDB facts used in a SLD-derivation of A wrt DDB, and construct that as an abductive explanation for A wrt IDB_G.

All solutions translate a view update request into a **transaction combining insertions and deletions of base relations** for satisfying the request [43]. Furthermore, a stratifiable (definite) deductive database can be considered as a knowledge base, and thus the rationality postulates and insertion algorithm from the previous section can be applied for solving view update requests in deductive databases.

Example 2. Consider a definite deductive database DDB as follows:

$$\begin{array}{lll} IDB: p \leftarrow a \wedge e & EDB: a \leftarrow & IC: \leftarrow b \\ \quad q \leftarrow a \wedge f & \quad e \leftarrow & \\ \quad p \leftarrow b \wedge f & \quad f \leftarrow & \\ \quad q \leftarrow b \wedge e & & \\ \quad p \leftarrow q & & \\ \quad q \leftarrow a & & \end{array}$$

Suppose We want to insert p. First, we need to check consistency with IC and afterwards, we have to find Δ_i and Δ_j via tree deduction.



It is easy to conclude which branches are consistent wrt IC (indicated in the depicted tree by the symbol ■). For the next step, we need to find minimal accommodate (positive literal) and denial literal (negative literal) with wrt to p. The subgoals of the tree are $\leftarrow a, e$ and $\leftarrow a, f$, which are minimal tree deductions of only facts. Clearly, $\Delta_i = \{a, e, f\}$ and $\Delta_j = \{b\}$ with respect to IC, are the only locally minimal abductive explanations for p wrt IDB_G, but they are not locally minimal explanations.

An algorithm for view update, based on the general revision algorithm (cf. Algorithm 1) in Section 2.1 and abductive explanation. There, given a view atom to be updated, set of all explanations for that atom has to be generated through a complete SLD-tree and a hitting set of these explanations is then update from the EDB. It was shown that this algorithm is rational. In this paper, we present a radically different approach that runs on polynomial space. The generation of hitting set is carried out through a hyper tableaux calculus (bottom-up) for deletion process and magic set(top-down) for insertion are focussed on the goal.

An algorithm for view updating can be developed based on the general revision algorithm and the generation of abductive explanations as proposed by Algorithm 1 and 2. For processing a given view update request, a set of all explanations for that atom has to be generated through a complete SLD-tree. The resulting hitting set of these explanations is then a base update of the EDB satsifying the view update request. In [12], it has been shown that this algorithm is rational. In this paper, we present a different approach which is also rational but even runs on polynomial space. The generation of a hitting set is carried out through a hyper tableaux calculus (bottom-up) for implementing the deletion process as well as through the magic sets approach (top-down) for performing insertions focussed on the particular goal given.

3.2 View update method

View updating [6] aims at determining one or more base relation updates such that all given update requests with respect to derived relations are satisfied after the base updates have been successfully applied.

Definition 7 (View update). *Let $DDB = \langle IDB, EDB, IC\rangle$ be a stratifiable (definite) deductive database $DDB(D)$. A VU request ν_D is a pair $\langle \nu_D^+, \nu_D^-\rangle$ where ν_D^+ and ν_D^- are sets of ground atoms representing the facts to be inserted into D or deleted from D, resp., such that $pred(\nu_D^+ \cup \nu_D^-) \subseteq pred(IDB)$, $\nu_D^+ \cap \nu_D^- = \emptyset$, $\nu_D^+ \cap PM_D = \emptyset$ and $\nu_D^- \subseteq PM_D$.*

Note that we consider again true view updates only, i.e., ground atoms which are presently not derivable for atoms to be inserted, or are derivable for atoms to be deleted, respectively. A method for view updating determines sets of alternative updates satisfying a given request. A set of updates leaving the given database consistent after its execution is called *VU realization*.

Definition 8 (Induced update). *Let $DDB = \langle IDB, EDB, IC\rangle$ be a stratifiable (definite) deductive database and $DDB = \nu_D$ a VU request. A VU realization is a base update u_D which leads to an induced update $u_{D\rightarrow D'}$ from D to D' such that $\nu_D^+ \subseteq PM_{D'}$ and $\nu_D^- \cap PM_{D'} = \emptyset$.*

There may be infinitely many realizations and even realizations of infinite size which satisfy a given VU request. A breadth-first search (BFS) is employed for determining a set of minimal realizations $\tau_D = \{u_D^1, \ldots, u_D^i\}$. Any u_D^i is minimal in the sense that none of its updates can be removed without losing the property of being a realization for ν_D.

Top-down computation: Given a VU request ν_{DDB}, view updating methods usually determine further VU requests in order to find relevant base updates. Similar to delta relations for UP we will use the notion VU relation to access individual view updates with respect to the relations of our system. For each relation $p \in pred(IDB \cup EDB)$ we use the VU relation $\nabla_p^+(\boldsymbol{x})$ for tuples to be inserted into DDB and $\nabla_p^-(\boldsymbol{x})$ for tuples to be deleted from DDB. The initial set of delta facts resulting from a given VU request is again represented by so-called *VU seeds*.

Definition 9 (View update seeds). *Let $DDB(D)$ be a stratifiable (definite) deductive database and $\nu_{DDB} = \langle \nu_D^+, \nu_D^-\rangle$ a VU request. The set of VU seeds $vu_seeds(\nu_D)$ with respect to ν_D is defined as follows:*

$$vu_seeds(\nu_D) := \left\{\nabla_p^\pi(c_1, \ldots, c_n) \mid p(c_1, \ldots, c_n) \in \nu_D^\pi \text{ and } \pi \in \{+, -\}\right\}.$$

Definition 10 (View update rules). *Let IDB be a normalized stratifiable (definite) deductive rule set. The set of VU rules for true view updates is denoted IDB^∇ and is defined as the smallest set satisfying the following conditions:*

1. *For each rule of the form* $p(\boldsymbol{x}) \leftarrow q(\boldsymbol{y}) \wedge r(\boldsymbol{z}) \in IDB$ *with* $vars(p(\boldsymbol{x})) = (vars(q(\boldsymbol{y})) \cup vars(r(\boldsymbol{z})))$ *the following three VU rules are in* IDB^{∇}*:*

$$\nabla_p^+(\boldsymbol{x}) \wedge \neg q(\boldsymbol{y}) \rightarrow \nabla_q^+(\boldsymbol{y}) \quad \nabla_p^-(\boldsymbol{x}) \rightarrow \nabla_q^-(\boldsymbol{y}) \vee \nabla_r^-(\boldsymbol{z})$$
$$\nabla_p^+(\boldsymbol{x}) \wedge \neg r(\boldsymbol{z}) \rightarrow \nabla_r^+(\boldsymbol{z})$$

2. *For each rule of the form* $p(\boldsymbol{x}) \leftarrow q(\boldsymbol{x}) \wedge \neg r(\boldsymbol{x}) \in IDB$ *the following three VU rules are in* IDB^{∇}*:*

$$\nabla_p^+(\boldsymbol{x}) \wedge \neg q(\boldsymbol{x}) \rightarrow \nabla_q^+(\boldsymbol{x}) \quad \nabla_p^-(\boldsymbol{x}) \rightarrow \nabla_q^-(\boldsymbol{x}) \vee \nabla_r^+(\boldsymbol{x})$$
$$\nabla_p^+(\boldsymbol{x}) \wedge r(\boldsymbol{x}) \rightarrow \nabla_r^-(\boldsymbol{x})$$

3. *For each two rules of the form* $p(\boldsymbol{x}) \leftarrow q(\boldsymbol{x})$ *and* $p(\boldsymbol{x}) \leftarrow r(\boldsymbol{x})$ *the following three VU rules are in* IDB^{∇}*:*

$$\nabla_p^-(\boldsymbol{x}) \wedge q(\boldsymbol{x}) \rightarrow \nabla_q^-(\boldsymbol{x}) \quad \nabla_p^+(\boldsymbol{x}) \rightarrow \nabla_q^+(\boldsymbol{x}) \vee \nabla_r^+(\boldsymbol{x})$$
$$\nabla_p^-(\boldsymbol{x}) \wedge r(\boldsymbol{x}) \rightarrow \nabla_r^-(\boldsymbol{x})$$

4. a) *For each relation* p *defined by a single rule* $p(\boldsymbol{x}) \leftarrow q(\boldsymbol{y}) \in IDB$ *with* $vars(p(\boldsymbol{x})) = vars(q(\boldsymbol{y}))$ *the following two VU rules are in* IDB^{∇}*:*

$$\nabla_p^+(\boldsymbol{x}) \rightarrow \nabla_q^+(\boldsymbol{y}) \quad \nabla_p^-(\boldsymbol{x}) \rightarrow \nabla_q^-(\boldsymbol{y})$$

 b) *For each relation* p *defined by a single rule* $p \leftarrow \neg q \in IDB$ *the following two VU rules are in* IDB^{∇}*:*

$$\nabla_p^+ \rightarrow \nabla_q^- \quad \nabla_p^- \rightarrow \nabla_q^+$$

5. *Assume without loss of generality that each projection rule in* IDB *is of the form* $p(\boldsymbol{x}) \leftarrow q(\boldsymbol{x}, Y) \in IDB$ *with* $Y \notin vars(p(\boldsymbol{x}))$*. Then the following two VU rules*

$$\nabla^- p_(\boldsymbol{x}) \wedge q(\boldsymbol{x}, Y) \rightarrow \nabla_q^-(\boldsymbol{x}, Y)$$
$$\nabla_p^+(\boldsymbol{x}) \rightarrow \nabla_q^+(\boldsymbol{x}, c_1) \vee \ldots \vee \nabla_q^+(\boldsymbol{x}, c_n) \vee \nabla_q^+(\boldsymbol{x}, c^{new})$$

 are in IDB^{∇} *where all* c_i *are constants from the Herbrand universe* $\mathcal{U}_{DDB}$ *of* DDB *and* c^{new} *is a new constant, i.e.,* $c^{new} \notin \mathcal{U}_{DDB}$*.*

Theorem 2. *Let* $DDB = \langle IDB, EDB, IC \rangle$ *be a stratifiable (definite)deductive database(D),* ν_D *a view update request and* $\tau_D = \{u_D^1, \ldots, u_D^n\}$ *the corresponding set of minimal realizations. Let* $D^{\nabla} = \langle EDB \cup vu_seeds(\nu_D), IDB \cup IDB^{\nabla} \rangle$ *be the transformed deductive database of* D*. Then the VU relations in* PM_D^{∇} *with respect to base relations of* D *correctly represent all direct consequences of* ν_D*. That is, for each realization* $u_D^i = \langle u_D^{i^+}, u_D^{i^-} \rangle \in \tau_D$ *the following condition holds:*

$$\exists p(\boldsymbol{t}) \in u_D^{i^+} : \nabla_p^+(\boldsymbol{t}) \in MS_D^{\nabla} \vee \exists p(\boldsymbol{t}) \in u_D^{i^-} : \nabla_p^-(\boldsymbol{t}) \in MS_D^{\nabla}.$$

Bottom-up computation: In [4,12] a variant of clausal normal form tableaux called "hyper tableaux" is introduced. Since the hyper tableaux calculus constitutes the basis for our view update algorithm, *Clauses,* i.e., multisets of literals, are usually written as the disjunction $A_1 \vee A_2 \vee \cdots \vee A_m \vee$ not $B_1 \vee$ not $B_2 \cdots \vee$ not B_n $(M \geq 0, n \geq 0)$. The literals $A_1, A_2, \ldots A_m$ (resp. $B_1, B_2, \ldots, B_n$) are called the *head (*resp. *body)* of a clause. With $\overline{L}$ we denote the complement of a literal L. Two literals L and K are complementary if $\overline{L} = K$.

From now on D always denotes a finite ground clause set, also called *database,* and Σ denotes its signature, i.e., the set of all predicate symbols occurring in it. We consider finite ordered trees T where the nodes, except the root node, are labeled with literals. In the following we will represent a branch b in T by the sequence $b = L_1, L_2, \ldots, L_n$ $(n \geq 0)$ of its literal labels, where L_1 labels an immediate successor of the root node, and L_n labels the leaf of b. The branch b is called *regular* iff $L_i \neq L_j$ for $1 \leq i, j \leq n$ and $i \neq j$, otherwise it is called *irregular.* The tree T is *regular* iff every of its branches is regular, otherwise it is *irregular.* The set of *branch literals* of b is $lit(b) = \{L_1, L_2, \ldots, L_n\}$. For brevity, we will write expressions like $A \in b$ instead of $A \in lit(b)$. In order to memorize the fact that a branch contains a contradiction, we allow to label a branch as either *open* or *closed.* A tableau is closed if each of its branches is closed, otherwise it is open.

Definition 11 (Hyper Tableau). *A literal set is called inconsistent* iff it contains a pair of complementary literals, otherwise it is called *consistent. Hyper tableaux* for D are inductively defined as follows:

Initialization step: The empty tree, consisting of the root node only, is a hyper tableau for D. Its single branch is marked as "open".

Hyper extension step: If (1) T is an open hyper tableau for D with open branch b, and (2) $C = A_1 \vee A_2 \vee \cdots \vee A_m \leftarrow B_1 \wedge B_2 \cdots \wedge B_n$ is a clause from D $(n \geq 0, m \geq 0)$, called *extending clause* in this context, and (3) $\{B_1, B_2, \ldots, B_n\} \subseteq b$ (equivalently, we say that C is *applicable to b)*then the tree T is a hyper tableau for D, where T is obtained from T by extension of b by C: replace b in T by the *new branches*

$$(b, A_1), (b, A_2), \ldots, (b, A_m), (b, \neg B_1), (b, \neg B_2), \ldots, (b, \neg B_n)$$

and then mark every inconsistent new branch as "closed", and the other new branches as "open".

The applicability condition of an extension expresses that all body literals have to be satisfied by the branch to be extended. From now on, we consider only regular hyper tableaux. This restriction guarantees that for finite clause sets no branch can be extended infinitely often. Hence, in particular, no open finished branch can be extended any further. This fact will be made use of below occasionally. Notice as an immediate consequence of the above definition that open branches never contain negative literals.

4 View update algorithm

The key idea of the algorithm presented in this paper is to transform the given database along with the view update request into a disjunctive logic program and apply known disjunctive techniques to solve the original view update problem. The intuition behind the transformation is to obtain a disjunctive logic program in such a way that each (minimal) model of this transformed program represent a way to update the given view atom. We present two variants of our algorithm. The one that is discussed in this section employs a trivial transformation procedure but has to look for minimal models; and another performs a costly transformation, but dispenses with the requirement of computing the minimal models.

4.1 Minimality test

We start presenting an algorithm for stratifiable (definite) deductive databases by first defining precisely how the given database is transformed into a disjunctive logic program for the view deletion process [12] (successful branch - see in [18] via Hyper Tableau).

Definition 12 (IDB Transformation). *Given an IDB and a set of ground atoms S, the transformation of IDB wrt S is obtained by translating each clause $C \in IDB$ as follows: Every atom A in the body (resp. head) of C that is also in S is moved to the head (resp. body) as $\neg A$.*

Note 1. If IDB is a stratifiable deductive database then the transformation introduced above is not necessary.

Definition 13 (IDB^* Transformation). *Let $IDB \cup EDB$ be a given database. Let $S_0 = EDB \cup \{A \mid A$ is a ground IDB atom$\}$. Then, IDB^* is defined as the transformation of IDB wrt S_0.*

Note 2. Note that IDB^* is in general a disjunctive logic program. The negative literals ($\neg A$) appearing in the clauses are intuitively interpreted as deletion of the corresponding atom (A) from the database. Technically, a literal $\neg A$ is to be read as a *positive* atom, by taking the $\neg$-sign as part of the predicate symbol. To be more precise, we treat $\neg A$ as an atom wrt IDB^*, but as a negative literal wrt IDB.

Note that there are no facts in IDB^*. So when we add a delete request such as $\neg A$ to this, the added request is the only fact and any bottom-up reasoning strategy is fully focused on the goal (here the delete request)

Definition 14 (Update Tableaux Hitting Set). *An update tableau for a database $IDB \cup EDB$ and delete request $\neg A$ is a hyper tableau T for $IDB^* \cup \{\neg A \leftarrow\}$ such that every open branch is finished. For every open finished branch b in T we define the hitting set (of b in T)* as $HS(b) = \{A \in EDB | \neg A \in b\}$.

Definition 15 (Minimality test). *Let T be an update tableau for $IDB \cup EDB$ and delete request $\neg A$. We say that open finished branch b in T satisfies the strong minimality test* iff $\forall s \in HS(b) : IDB \cup EDB \backslash HS(b) \cup \{s\} \vdash A$.

Definition 16 (Update Tableau satisfying strong minimality). *An update tableau for given $IDB \cup EDB$ and delete request $\neg A$ is transformed into an update tableau satisfying strong minimality by marking every open finished branch as closed which does not satisfy strong minimality.*

The next step is to consider the view insertion process [6] (unsuccessful branch - see [18]).

Definition 17 (IDB^* Transformation). *Let $IDB \cup EDB$ be a given database. Let $S_1 = EDB \cup \{A \mid A$ is a ground IDB* atom$\}$. *Then, IDB^* is defined as the transformation of IDB wrt S_1.*

Note 3. Note that IDB is in general a (stratifiable) disjunctive logic program. The positive literals (A) appearing in the clauses are intuitively interpreted as an insertion of the corresponding atom (A) from the database.

Definition 18 (Update magic Hitting Set). *An update magic set rule for a database $IDB \cup EDB$ and insertion request A is a magic set rule M for $IDB^* \cup \{A \leftarrow\}$ such that every close branch is finished. For every close finished branch b in M we define the magic set rule (of b in M)* as $HS(b) = \{A \in EDB | A \in b\}$.

Definition 19 (Minimality test). *Let M be an update magic set rule for $IDB \cup EDB$ and insert request A. We say that close finished branch b in M satisfies the strong minimality test* iff $\forall s \in HS(b) : IDB \cup EDB \backslash HS(b) \cup \{s\} \vdash \neg A$.

Definition 20 (Update magic set rule satisfying strong minimality). *An update magic set rule for given $IDB \cup EDB$ and insert request A is transformed into an update magic set rule satisfying strong minimality by marking every close finished branch as open which does not satisfy strong minimality.*

Algorithm 3	View updating Algorithm based on minimality test
Input :	A definite deductive database $DDB = IDB \cup EDB \cup IC$
Output:	A new database $IDB \cup EDB' \cup IC$
begin	
1.	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	While $(V \neq \emptyset)$
2.	For every successful branch i:construct $\Delta_i = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch i.
	Construct a branch i of an update tableau satisfying minimality
	for $IDB \cup EDB$ and delete request $\neg A$.
	Produce $IDB \cup EDB \backslash HS(i)$ as a result
3.	For every unsuccessful branch j:construct $\Delta_j = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch j.
	Construct a branch j of an update magic set rule satisfying minimality
	for $IDB \cup EDB$ and insert request A.
	Produce $IDB \cup EDB \backslash HS(j)$ as a result
	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	return
5.	Produce DDB as the result.
end.	

Algorithm 4	View updating Algorithm based on minimality test
Input :	A stratifiable deductive database $DDB = IDB \cup EDB \cup IC$
Output:	A new database $IDB \cup EDB' \cup IC$
begin	
1.	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	While $(V \neq \emptyset)$
2.	For every successful branch i:construct $\Delta_i = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch i.
	Construct a branch i of an update tableau satisfying minimality
	for $IDB \cup EDB$ and delete request A.
	Produce $IDB \cup EDB \backslash HS(i)$ as a result
3.	For every unsuccessful branch j:construct $\Delta_j = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch j.
	Construct a branch j of an update magic set rule satisfying minimality
	for $IDB \cup EDB$ and insert request A.
	Produce $IDB \cup EDB \backslash HS(j)$ as a result
	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	return
5.	Produce DDB as the result.
end.	

Lemma 1. *The strong minimality test and the groundedness test are equivalent.*

This means that every minimal model (minimal wrt the base atoms) of $IDB^* \cup \{\neg A\}$ provides a minimal hitting set for deleting the ground view atom A. Similarly, $IDB^* \cup \{A\}$ provides a minimal hitting set for inserting the ground view atom A. Now we are in a position to formally present our algorithm. Given a database and a view atom to be updated, we first transform the database into a definite disjunctive logic program and use hyper tableaux calculus to generate models of this transformed program for deletion of an atom. Second, magic sets transformed rules are used is used to generate models of this transformed program for determining an induced insertion of an atom. Models that do not represent rational update are filtered out using the strong minimality test. This is formalized in Algorithm 3. The procedure for stratifiable deductive databases is presented in Algorithm 4.

To show the rationality of this approach, we study how this is related to the previous approach presented in the last section, i.e., generating explanations and computing hitting sets of these explanations. To better understand the relationship it is imperative to study where the explanations are in the hyper tableau approach and magic set rules. We first define the notion of an EDB -cut and then view update seeds.

Definition 21 (EDB-Cut). *Let T be update tableau with open branches $b_1, b_2, \ldots, b_n$. A set $S = \{A_1, A_2, \ldots, A_n\} \subseteq EDB$ is said to be EDB-cut of T iff $\neg A_i \in b_i$ ($A_i \in b_i$), for $1 \leq i \leq n$.*

Definition 22 (EDB seeds). *Let M be an update seeds with close branches $b_1, b_2, \ldots, b_n$. A set $S = \{A_1, A_2, \ldots, A_n\} \subseteq EDB$ is said to be a EDB-seeds of M iff EDB seeds $vu_seeds(\nu_D)$ with respect to ν_D is defined as follows:*

$$vu_seeds(\nu_D) := \left\{\nabla_p^{\pi}(c_1, \ldots, c_n) | p(c_1, \ldots, c_n) \in \nu_D^{\pi} \text{ and } \pi \in \{+, -\}\right\}.$$

Lemma 2. *Let T be an update tableau for $IDB \cup EDB$ and update request A. Similarly, for M be an update magic set rule. Let S be the set of all EDB-closed minimal abductive explanations for A wrt. IDB. Let S' be the set of all EDB-cuts of T and EDB-seeds of M . Then the following hold*

- $S \subseteq S'$.
- $\forall \Delta' \in S' : \exists \Delta \in S s.t. \Delta \subseteq \Delta'$.

The above lemma precisely characterizes what explanations are generated by an update tableau. It is obvious then that a branch cuts through all the explanations and constitutes a hitting set for all the generated explanations. This is formalized below.

Lemma 3. *Let S and S' be sets of sets s.t. $S \subseteq S'$ and every member of $S' \backslash S$ contains an element of S. Then, a set H is a minimal hitting set for S iff it is a minimal hitting set for S'.*

Lemma 4. *Let T be an update tableau for $IDB \cup EDB$ and update request A that satisfies the strong minimality test. Similarly, for M be an update magic set rule. Then, for every open (close) finished branch b in T, $HS(b)$ (M, $HS(b)$) is a minimal hitting set of all the abductive explanations of A.*

So, Algorithms 3 and 4 generate a minimal hitting set (in polynomial space) of all EDB-closed locally minimal abductive explanations of the view atom to be deleted. From the belief dynamics results recalled in section 3, it immediately follows that Algorithms 5 and 6 are rational, and satisfy the strong relevance postulate (KB-7.1).

Theorem 3. *Algorithms 3 and 4 are rational, in the sense that they satisfy all the rationality postulates (KB*1)-(KB*6) and the strong relevance postulate (KB*7.1). Further, any update that satisfies these postulates can be computed by these algorithms.*

4.2 Materialized view

In many cases, the view to be updates is materialized, i.e., the least Herbrand Model is computed and kept, for efficient query answering. In such a situation, rational hitting sets can be computed without performing any minimality test. The idea is to transform the given IDB wrt the materialized view.

Definition 23 (IDB^+ Transformation). *Let $IDB \cup EDB$ be a given database. Let S be the Least Herbrand Model of this database. Then, IDB^+ is defined as the transformation of IDB wrt S.*

Note 4. If IDB is a stratifiable deductive database then the transformation introduced above is not necessary.

Definition 24 (Update Tableau based on Materialized view). *An update tableau based on materialized view for a database $IDB \cup EDB$ and delete request $\neg A$ is a hyper tableau T for $IDB^+ \cup \{\neg A \leftarrow\}$ such that every open branch is finished.*

Definition 25 (IDB^- Transformation). *Let $IDB \cup EDB$ be a given database. Let S_1 be the Least Herbrand Model of this database. Then, IDB^- is defined as the transformation of IDB wrt S_1.*

Definition 26 (Update magic set rule based on Materialized view). *An update magic set rule based on materialized view for a database $IDB \cup EDB$ and insert request A is a magic set M for $IDB^+ \cup \{A \leftarrow\}$ such that every close branch is finished.*

Now the claim is that every model of $IDB^+ \cup \{\neg A \leftarrow\}$ ($A \leftarrow$) constitutes a rational hitting set for the deletion and insertion of the ground view atom A. So, the algorithm works as follows:

Algorithm 5	View update algorithm based on Materialized view
Input :	A definite deductive database $DDB = IDB \cup EDB \cup IC$
Output:	A new database $IDB \cup EDB' \cup IC$
begin	
1.	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	While $(V \neq \emptyset)$
2.	For every successful branch i:construct $\Delta_i = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch i.
	Construct a branch i of an update tableau based on view
	for $IDB \cup EDB$ and delete request $\neg A$.
	Produce $IDB \cup EDB \backslash HS(i)$ as a result
3.	For every unsuccessful branch j:construct $\Delta_j = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch j.
	Construct a branch j of an update magic set rule based on view
	for $IDB \cup EDB$ and insert request A.
	Produce $IDB \cup EDB \backslash HS(j)$ as a result
	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	return
5.	Produce DDB as the result.
end.	

Algorithm 6	View update algorithm based on Materialized view
Input :	A stratifiable deductive database $DDB = IDB \cup EDB \cup IC$
Output:	A new database $IDB \cup EDB' \cup IC$
begin	
1.	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	While $(V \neq \emptyset)$
2.	For every successful branch i:construct $\Delta_i = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch i.
	Construct a branch i of an update tableau satisfying based on view
	for $IDB \cup EDB$ and delete request A.
	Produce $IDB \cup EDB \backslash HS(i)$ as a result
3.	For every unsuccessful branch j:construct $\Delta_j = \{D \mid D \in EDB\}$
	and D is used as an input clause in branch j.
	Construct a branch j of an update magic set rule based on view
	for $IDB \cup EDB$ and insert request A.
	Produce $IDB \cup EDB \backslash HS(j)$ as a result
	Let $V := \{c \in IC \mid IDB \cup IC$ inconsistent with $\mathcal{A}$ wrt $c\}$
	return
5.	Produce DDB as the result.
end.	

Given a database and a view update request, we first transform the database wrt its Least Herbrand Model (computation of the Least Herbrand Model can

be done as a offline preprocessing step. Note that it serves as materialized view for efficient query answering). Then the hyper tableaux calculus (magic set rule) is used to compute models of this transformed program. Each model represents a rational way of accomplishing the given view update request. This is formalized in Algorithms 5 and 6.

This approach for view update may not satisfy (KB*7.1) in general. But, as shown in the sequel, conformation to (KB*6.3) is guaranteed and thus this approach results in rational update.

Lemma 5. *Let T be an update tableau based on a materialized view for $IDB \cup EDB$ and delete request $\neg A$ (A), Similarly, let M be an update magic set rule. Let S be the set of all EDB-closed locally minimal abductive explanations for A wrt IDB. Let S' be the set of all EDB-cuts of T and EDB-seeds of M. Then, the following hold:*

- $S \subseteq S'$.
- $\forall \Delta' \in S' : \exists \Delta \in S$ *s.t.* $\Delta \subseteq \Delta'$.
- $\forall \Delta' \in S' : \Delta' \subseteq \bigcup S$.

Lemma 6. *Let S and S' be sets of sets s.t. $S \in S'$ and for every member X of $S' \backslash S$: X contains a member of S and X is contained in $\bigcup S$. Then, a set H is a hitting set for S iff it is a hitting set for S'.*

Lemma 7. *Let T and M as in Lemma 5. Then $HS(b)$ is a rational hitting set for A, for every open finished branch b in T (close finished branch b in M).*

Theorem 4. *Algorithms 5 and 6 are rational, in the sense that they satisfy all the rationality postulates (KB*1) to (KB*6) and (KB*7.3).*

5 Related Works

We begin by recalling previous work on view deletion. Chandrabose [10,11] and Delhibabu [18,19], defines a contraction and revision operator in view deletion with respect to a set of formulae or sentences using Hansson's [24] belief change. Similar to our approach, he focused on set of formulae or sentences in knowledge base revision for view update wrt. insertion and deletion and formulae are considered at the same level. Chandrabose proposed different ways to change knowledge base via only database deletion, devising particular postulate which is shown to be necessary and sufficient for such an update process.

Our Horn knowledge base consists of two parts, immutable part and updatable part, but our focus is on minimal change computations. There is more related works on that topic. Eiter [20], Langlois[31], and Delgrande [17] are focusing on Horn revision with different perspectives like prime implication, logical closure and belief level. Segerberg [53] defined a new modeling technique for belief revision in terms of irrevocability on prioritized revision. Hansson [24], constructed five types of non-prioritized belief revision. Makinson [40] developed dialogue form of revision AGM. Papini[48] defined a new version of knowledge

base revision. In this paper, we considered the immutable part as a Horn clause and the updatable part as an atom (literal).

Hansson's[24] kernel change is related to abductive method. Aliseda's [2] book on abductive reasoning is one of the motivation keys. Christiansen's [15] work on dynamics of abductive logic grammars exactly fits our minimal change (insertion and deletion). Wrobel's [56] definition of first order theory revision was helpful to frame our algorithm.

On other hand, we are dealing with view update problem. Keller's [27] thesis is motivation the view update problem. There is a lot of papers on the view update problem (for example, the recent survey paper on view updating by Chen and Liao [14] and the survey paper on view updating algorithms by Mayol and Teniente [41] and current survey paper on view selection ([49,3,5]). More similar to our work is the paper presented by Bessant et al. [7], which introduces a local search-based heuristic technique that empirically proves to be often viable, even in the context of very large propositional applications. Laurent et al.[32], considers updates in a deductive database in which every insertion or deletion of a fact can be performed in a deterministic way.

Furthermore, and at a first sight more related to our work, some work has been done on ontology systems and description logics (Qi and Yang [50],Kogalovsky [28] and Zang [58]). In Fuzzy related work ([57,35,36,55,47]) also in the current attenuation of database people.

The significance of our work can be summarized in the following:

- We have defined a new kind of kernel operator on knowledge bases and obtained an axiomatic characterization for it. This operator of change is based on α consistent-remainder set. Thus, we have presented a way to construct a kernel operator without the need to make use of the generalized Levi's identity nor of a previously defined revision operator.
- We have defined a new way of insertion and deletion of an atom(literals) as per norm of principle of minimal change.
- We have proposed a new generalized revision algorithm for knowledge base dynamics, interesting connections with kernel change and abduction procedure.
- We have designed a new view update algorithm for stratifiable DDB, using an axiomatic method based on Hyper tableaux and magic sets.

6 Conclusion and remarks

The main contribution of this research is to provide a link between theory of belief dynamics and concrete applications such as view updates in databases. We argued for generalization of belief dynamics theory in two respects: to handle certain part of knowledge as immutable; and dropping the requirement that belief state be deductively closed. The intended generalization was achieved by introducing the concept of knowledge base dynamics and generalized revision for the same. Further, we also studied the relationship between knowledge base dynamics and abduction resulting in a generalized algorithm for revision

based on abductive procedures. We also successfully demonstrated how knowledge base dynamics can provide an axiomatic characterization for updating an atom(literals) to a stratifiable (definite) deductive database.

In bridging the gap between belief dynamics and view updates, we have observed that a balance has to be achieved between computational efficiency and rationality. While rationally attractive notions of generalized revision prove to be computationally inefficient, the rationality behind efficient algorithms based on incomplete trees is not clear at all. From the belief dynamics point of view, we may have to sacrifice some postulates, vacuity for example, to gain computational efficiency. Further weakening of relevance has to be explored, to provide declarative semantics for algorithms based on incomplete trees.

On the other hand, from the database side, we should explore various ways of optimizing the algorithms that would comply with the proposed declarative semantics. We believe that partial deduction and loop detection techniques, will play an important role in optimizing algorithms of the previous section. Note that, loop detection could be carried out during partial deduction, and complete SLD-trees can be effectively constructed wrt a partial deduction (with loop check) of a database, rather than wrt database itself. Moreover, we would anyway need a partial deduction for optimization of query evaluation.

We have presented two variants of an algorithm for update a view atom from a definite database. The key idea of this approach is to transform the given database into a disjunctive logic program in such a way that updates can be read off from the models of this transformed program. One variant based on materialized views is of polynomial time complexity. Moreover, we have also shown that this algorithm is rational in the sense that it satisfies the rationality postulates that are justified from philosophical angle.

In the second variant, where materialized view is used for the transformation, after generating a hitting set and removing corresponding EDB atoms, we easily move to the new materialized view. An obvious way is to recompute the view from scratch using the new EDB (i.e., compute the Least Herbrand Model of the new updated database from scratch) but it is certainly interesting to look for more efficient methods.

Though we have discussed only about view updates, we believe that knowledge base dynamics can also be applied to other applications such as view maintenance, diagnosis, and we plan to explore it further (see works [9] and [8]). It would also be interesting to study how results using soft stratification [6] with belief dynamics, especially the relational approach, could be applied in real world problems. Still, a lot of developments are possible, for improving existing operators or for defining new classes of change operators. As immediate extension, question raises: is there any *real life application for AGM in 25 year theory?* [22]. The revision and update are more challenging in logical view update problem(database theory), so we can extend the theory to combine results similar to Konieczny's [29] and Nayak's [45].

Appendix

Proof of Theorem 1. Follows from Algorithm 1 and 2. ■

Proof of Theorem 2. Follows from the result of [6] ■

Proof of Lemma 1. Follows from the result of [12]

Proof of Lemma 2 and 5.

1. Consider a $\Delta(\Delta \in \Delta_i \cup \Delta_j) \in S$. We need to show that Δ is generated by algorithm 3 at step 2. It is clear that there exists a A-kernel X of DDB_G s.t. $X \cap EDB = \Delta_j$ and $X \cup EDB = \Delta_i$. Since $X \vdash A$, there must exist a successful derivation for A using only the elements of X as input clauses and similarly $X \nvdash A$. Consequently Δ must have been constructed at step 2.
2. Consider a $\Delta'((\Delta' \in \Delta_i \cup \Delta_j) \in S'$. Let Δ' be constructed from a successful(unsuccessful) branch i via $\Delta_i(\Delta_j)$. Let X be the set of all input clauses used in the refutation i. Clearly $X \vdash A(X \nvdash A)$. Further, there exists a minimal (wrt set-inclusion) subset Y of X that derives A (i.e., no proper subset of Y derives A). Let $\Delta = Y \cap EDB$ ($Y \cup EDB$). Since IDB does not(does) have any unit clauses, Y must contain some EDB facts, and so Δ is not empty (empty) and obviously $\Delta \subseteq \Delta'$. But, Y need not (need) be a A-kernel for IDB_G since Y is not ground in general. But it stands for several A-kernels with the same (different) EDB facts Δ in them. Thus, from lemma 1, Δ is a DDB-closed locally minimal abductive explanation for A wrt IDB_G and is contained in Δ'.

Proof of Lemma 3 and 6.

1. (**Only if part**) Suppose H is a minimal hitting set for S. Since $S \subseteq S'$, it follows that $H \subseteq \bigcup S'$. Further, H hits every element of S' , which is evident from the fact that every element of S' contains an element of S. Hence H is a hitting set for S' . By the same arguments, it is not difficult to see that H is minimal for S' too.

 (**If part**) Given that H is a minimal hitting set for S' , we have to show that it is a minimal hitting set for S too. Assume that there is an element $E \in H$ that is not in $\bigcup S$. This means that E is selected from some $Y \in S' \backslash S$. But Y contains an element of S, say X. Since X is also a member of S' , one member of X must appear in H. This implies that two elements have been selected from Y and hence H is not minimal. This is a contradiction and hence $H \subseteq \bigcup S$. Since $S \subseteq S'$, it is clear that H hits every element in S, and so H is a hitting set for S. It remains to be shown that H is minimal. Assume the contrary, that a proper subset H' of H is a hitting set for S. Then from the proof of the only if part, it follows that H' is a hitting set for S' too, and contradicts the fact that H is a minimal hitting set for S' . Hence, H must be a minimal hitting set for S.

2. (**If part**) Given that H is a hitting set for S' , we have to show that it is a hitting set for S too. First of all, observe that $\bigcup S = \bigcup S'$, and so $H \subseteq \bigcup S$. Moreover, by definition, for every non-empty member X of S' , $H \cap X$ is not empty. Since $S \subseteq S'$, it follows that H is a hitting set for S too.

 (**Only if part**) Suppose H is a hitting set for S. As observed above, $H \subseteq \bigcup S'$. By definition, for every non-empty member $X \in S$, $X \cap H$ is not empty. Since every member of S' contains a member of S, it is clear that H hits every member of S' , and hence a hitting set for S' . ■

Proof of Lemma 4 and 7. Follows from the lemma 2,3 (minimal test) and 5,6 (materialized view) of [6] ■

Proof of Theorem 3. Follows from Lemma 4 and Theorem 1. ■

Proof of Theorem 4. Follows from Lemma 7 and Theorem 3. ■

Acknowledgement

The author acknowledges the support of RWTH Aachen, where he is visiting scholar with an Erasmus Mundus External Cooperation Window India4EU by the European Commission when the paper was written. I would like to thanks Chandrabose Aravindan and Gerhard Lakemeyer both my Indian and Germany PhD supervisor, give encourage to write the paper.

References

1. Alchourron, C.E., et al.(1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic 50*, 510 - 530.
2. Aliseda, A. (2006). *Abductive Resoning Logic Investigations into Discovery and Explanation*. Springer book series Vol. 330.
3. Amirkhani, H. & Rahmati. M. (2014). Agreement/disagreement based crowd labeling, *Applied Intelligence*, Accepted.
4. Baumgartner, P., et al. (1997). Semantically Guided Theorem Proving for Diagnosis Applications. *IJCAI* **1**, 460-465.
5. Baumeister, J., et al. (2011). KnowWE: a Semantic Wiki for knowledge engineering, *Applied Intelligence*, **35**(3), 323-344.
6. Behrend, A.,& Manthey,R. (2008). A Transformation-Based Approach to View Updating in Stratifiable Deductive Databases. *FoIKS*, 253-271.
7. Bessant, B., et al.(1998). Combining Nonmonotonic Reasoning and Belief Revision: A Practical Approach. *AIMSA*, 115-128.
8. Biskup, J. (2012). Inference-usability confinement by maintaining inference-proof views of an information system. *IJCSE*. **7** (1), 17-37.
9. Caroprese, L., et al.(2012). The View-Update Problem for Indefinite Databases. *JELIA*.
10. Chandrabose, A.,& Dung, P.M.(1994). Belief Dynamics, Abduction, and Database. *JELIA*, 66-85.

11. Chandrabose A.(1995), *Dynamics of Belief: Epistmology, Abduction and Database Update*. Phd Thesis, AIT.
12. Chandrabose. A., & Baumgartner. P. (1997). A Rational and Efficient Algorithm for View Deletion in Databases. *ILPS*, 165-179.
13. Calvanese, D., et al. (2012). View-based query answering in Description Logics Semantics and complexity. *J. Comput. Syst. Sci* **78**(1), 26-46.
14. Chen, H., & Liao, H. (2010). A Comparative Study of View Update Problem. *DSDE*, 83-89.
15. Christiansen, H., & Dahl,V. (2009). Abductive Logic Grammars. *WoLLIC*, 170-181.
16. Cong, G., et al. (2012). On the Complexity of View Update Analysis and Its Application to Annotation Propagation. *IEEE Trans. Knowl. Data Eng.***24**(3), 506-519.
17. Delgrande, J.P, & Peppas, P. (2011). Revising Horn Theories. *IJCAI*, 839-844.
18. Delhibabu, R., & Lakemeyer, G. (2013). A Rational and Efficient Algorithm for View Revision in Databases. *Applied Mathematics & Information Sciences*, **7**, pp: 843-856.
19. Delhibabu, R. (2014). An Abductive Framework for Knowledge Base Dynamics. *Applied Mathematics & Information Sciences* (accepted).
20. Eiter, T., & Makino,K. (2007). On computing all abductive explanations from a propositional Horn theory. *J. ACM* **54** (5).
21. Falappa, M.A., et al.(2012). Prioritized and Non-prioritized Multiple Change on Belief Bases. *J. Philosophical Logic* **41** (1), 77-113.
22. Fermé, E.L., & Hansson, S.O. (2011). AGM 25 Years - Twenty-Five Years of Research in Belief Change. *J. Philosophical Logic* **40** (2),295-331.
23. Hansson, S.O. (1991). Belief contraction without recovery. *Studia Logica* **50**(2), 251 Ű 260.
24. Hansson, S.O. (1997).*A Textbook of Belief Dynamics*. Kluwer Academic Publishers, Dordrecht.
25. Herzig, A. & Rifi,O. (1999). Propositional Belief Base Update and Minimal Change. *Artif. Intell.* **115**(1), 107-138.
26. Jackson, E. K. & Schulte, W. (2008). Model Generation for Horn Logic with Stratified Negation. FORTE.
27. Keller, A. (1985).*Updating Relational Databases Through Views*. Phd Thesis.
28. Kogalovsky, M.R. (2012). Ontology-based data access systems. *Programming and Computer Software* **38**(4), 167-182.
29. Konieczny, S. (2011). Dynamics of Beliefs. *SUM*, 61-74.
30. Lakemeyer, G. (1995). A Logical Account of Relevance. *IJCAI* (1), 853-861.
31. Langlois, M., et al. (2008). Horn Complements: Towards Horn-to-Horn Belief Revision. *AAAI*, 466-471.
32. Laurent, D., et al. (1998). Updating Intensional Predicates in Deductive Databases. *Data Knowl. Eng.* **26**(1), 37-70.
33. Liberatore, P. (1997). The Complexity of Belief Update (Extended in 2003). *IJCAI*(1), 68-73.
34. Liberatore, P., & Schaerf, M. (2004). The Compactness of Belief Revision and Update Operators. *Fundam. Inform.* **62**(3-4), 377-393.
35. Lin, J., et al. (2013). Storing and querying fuzzy XML data in relational databases, *Applied Intelligence*, **39**(2), 386-396.
36. Lin, J. & Ma. M. (2013). Formal transformation from fuzzy object-oriented databases to fuzzy XML, *Applied Intelligence*, **39**(3), 630-641.

37. Lobo, J., Minker, J., & Rajasekar, A. (1992). Foundations of Disjunctive Logic Programming. *MIT Press, Cambridge*.
38. Lobo, J., & Trajcevski, G. (1997). Minimal and Consistent Evolution of Knowledge Bases, *Journal of Applied Non-Classical Logics* **7**(1).
39. Lu, W. (1999). View Updates in Disjunctive Deductive Databases Based on SLD-Resolution. *KRDB*, 31-35.
40. Makinson, D. (1997). Screened Revision, *Theoria* **63**, 14-23.
41. Mayol, E., & Teniente, E. (1999). A Survey of Current Methods for Integrity Constraint Maintenance and View Updating. *ER (Workshops)*, 62-73.
42. Meyden, R. (1998). Logical Approaches to Incomplete Information: A Survey. *Logics for Databases and Information Systems*, 307-356.
43. Mota-Herranz, L., et al. (2000). Transaction Trees for Knowledge Revision, *FQAS*, 182-191.
44. Nayak, A., et al. (2006). Forgetting and Knowledge Update. *Australian Conference on Artificial Intelligence*, 131-140.
45. Nayak, A. (2011). Is Revision a Special Kind of Update? *Australasian Conference on Artificial Intelligence*, 432-441.
46. Nebel, B. (1998). How Hard is it to Revise a Belief Base? *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, 77-145.
47. Papadakis, N., et al. (2012). The ramification problem in temporal databases: a solution implemented in SQL, *Applied Intelligence*, **36**(4), 749-767.
48. Papini, O.(2000). Knowledge-base revision. *The Knowledge Engineering Review* **15**(4), 339 - 370.
49. Potter, s. (2012). Critical reasoning: AI for emergency response, *Applied Intelligence*, **37**(3), 337-356.
50. Qi, G., & Yang, F. (2008). A Survey of Revision Approaches in Description Logics. *Description Logics*.
51. Sakama, C., & Inoue, K. (2003). An abductive framework for computing knowledge base updates. *TPLP* **3**(6), 671-713.
52. Schulte, O. (1999). Minimal Belief Change and Pareto-Optimality. *Australian Joint Conference on Artificial Intelligence*, 144-155.
53. Segerberg, K. (1998). Irrevocable Belief Revision in Dynamic Doxastic Logic. *Notre Dame Journal of Formal Logic* **39**(3), 287-306.
54. Teniente, E., & Urpí, T. (2003). On the abductive or deductive nature of database schema validation and update processing problems. *TPLP* **3** (3), 287-327.
55. Wang, J., et al. (2012). On the combination of logical and probabilistic models for information analysis, *Applied Intelligence*, **36**(2), 472-497.
56. Wrobel, S. (1995). First order Theory Refinement. *IOS Frontier in AI and Application Series*.
57. Xu, C., et al. (2012). Efficient fuzzy ranking queries in uncertain databases, *Applied Intelligence*, **37**(1), 47-59.
58. Zhang, F, Z. & Ma. M. (2014). Representing and Reasoning About XML with Ontologies, *Applied Intelligence*, **40**(1), 74-106.